

CLAIMS

What is claimed is:

1. A computer-readable medium having computer-executable instructions for performing steps for backing up and restoring a data set of a node in a distributed system, comprising:
 - initiating a back up operation;
 - setting state data of the data set to indicate that the data set is in a backed up state; and
 - taking a snapshot of the contents of the data set including the state data to generate a backup copy, the contents of the data set including a plurality of changes each identified by a replication identification number of the node and a serial number;
 - resetting the state data of the data set to indicate that the data set is in normal operation after taking the snapshot.
2. A computer-readable medium as in claim 1, having further computer-executable instructions for performing the steps of:
 - restoring the data set using the backup copy;
 - detecting that the stat data of the restored data set indicates that the data set is in the backed up state;
 - in response to the detecting, changing the replication identification number of the node from an old value used before the backup operation to a new value.
3. A computer-readable medium as in claim 2, having further computer-executable instructions for performing the step of storing into the data set a lowest uncommitted serial

number that corresponds to a lowest one of the serial numbers of changes made to the data set that are not yet committed prior to taking the snapshot.

4. A computer-readable medium as in claim 3, having further computer-readable instructions for performing the step of requesting a second node in the distributed system to replicate changes that have the old value of the replication identification number of the node and serial numbers equal or higher than said lowest uncommitted serial number.

5. A computer-readable medium as in claim 2, having further computer-executable instructions for performing the step of storing in the data set, prior to taking the snapshot, a next serial number to be assigned to a new change to the data set.

6. A computer-readable medium as in claim 1, wherein the replication identification number of the node is a computed Globally Unique Identifier (GUID).

7. A computer-readable medium as in claim 1, wherein the distributed system is a directory service system and the data set of the node contains directory data.

8. A method for backing up and restoring contents of a data set of a node in a distributed system, comprising:

initiating a back up operation;

setting state data of the data set to indicate that the data set is in a backed up state; and

taking a snapshot of the contents of the data set including the state data to generate a backup copy, the contents of the data set including a plurality of changes each identified by a replication identification number of the node and a serial number;

resetting the state data of the data set to indicate that the data set is in normal operation after taking the snapshot.

9. A method as in claim 8, further including the steps of:

restoring the data set using the backup copy;

detecting that the state data of the restored data set indicates that the data set is in the backed up state;

in response to the detecting, changing the replication identification number of the node from an old value used before the backup operation to a new value.

10. A method as in claim 9, further including the step of storing into the data set a lowest uncommitted serial number that corresponds to a lowest one of serial numbers of changes made to the data set that are not yet committed prior to taking the snapshot.

11. A method as in claim 10, further including the step of requesting a second node in the distributed system to replicate changes that have the old value of the replication identification number of the node and serial numbers equal to or higher than said lowest uncommitted serial number.

12. A method as in claim 9, further including the step of storing in the data set, prior to taking the snapshot, a next serial number to be assigned to a new change to the data set.

13. A method as in claim 8, further including the step of setting the replication identification number of the node as a computed Globally Unique Identifier (GUID).

14. A method as in claim 8, wherein the distributed system is a directory service system and the data set of the node contains directory data.

15. A computer system for implementing a node in a distributed system, comprising:
a data manager for maintaining a data set containing a plurality of changes each identified by a replication identification number of the node and a serial number;
a replication server for replicating changes made to the data set to other nodes in the distributed system;
a backup server for generating a backup copy of the data set, the backup server being programmed to cause state data of the data set to be set to indicate that the data set is in a backed up state, take a snapshot of the contents of the data set including the state data to generate the backup copy, and cause the state data of the data set to be reset to indicate that the data set is in normal operation after taking the snapshot.

16. A computer system as in claim 15, wherein the replication server is programmed to detect, after the data set is restored from the backup copy, that the state data of the restored data

set indicates that the data set is in the backed up state, and change the replication identification number of the node from an old value used before the snapshot to a new value.

17. A computer system as in claim 16, wherein the replication server is further programmed to store into the data set a lowest uncommitted serial number that corresponds to a lowest one of serial numbers of changes made to the data set that are not yet committed prior to taking the snapshot.

18. A computer system as in claim 17, wherein the replication server is further programmed to request a second node in the distributed system to replicate changes that have the old value of the replication identification number of the node and serial numbers equal or higher than said lowest uncommitted serial number.

19. A computer system as in claim 16, wherein the replication server is further programmed to store in the data set, prior to taking the snapshot, a next serial number to be assigned to a new change to the data set.

20. A computer system as in claim 15, wherein the replication identification number of the node is a computed Globally Unique Identifier (GUID).

21. A computer system as in claim 15, wherein the distributed system is a directory service system and the data set of the node contains directory data.